

High-end server low-temperature cooling

by R. R. Schmidt
B. D. Notohardjono

The IBM S/390® G4 CMOS system, first shipped in 1997, was the first high-end system to use refrigeration. The decision to employ refrigeration cooling instead of other cooling options such as high-flow air cooling or various water-cooling schemes focused on the potential system performance improvement obtainable by lowering coolant temperatures using a refrigeration system. This paper reviews the historical background of refrigeration from its use in the early 1800s to its implementation in computer systems in the early 1990s. The advantages and disadvantages of using refrigeration in the cooling of computer systems are examined. The advantages have outweighed the disadvantages, leading to the first use by IBM of refrigeration in cooling the S/390 G4 server. The design of the refrigeration system for the S/390 G4 system is described in detail, and some of the key parametric studies that contributed to the final design are described.

Introduction

Moore's law, proposed by Gordon Moore in 1965 [1] and updated in 1975 [2], stated that semiconductor performance would double every 18 months. Throughout the last 25 years this law has stood the test of time, largely

because of device scaling. However, unless a new approach is identified for improving performance in complementary metal-oxide semiconductor (CMOS) devices, this law will cease to be valid. In the year 2000, scientists predicted [3] that improvements in performance would continue between twelve and fifteen years, depending on whether the current technology continues in use or transistors with new materials and new structures are developed. In any case, it is generally agreed that we are approaching a limit in the improvement in semiconductor performance.

Because of this limit, operation of CMOS semiconductor devices at low temperature to improve performance has been considered. (The term "low temperature" refers here to any temperatures lower than temperatures of air-cooled chips, which typically operate between 60 and 100°C.) The potential for low-temperature enhancement of CMOS performance has been recognized for some time, going back as far as the late 1960s and mid-1970s. The advantages of operating electronics at low temperatures, which include faster switching times of semiconductor devices and increased circuit speed due to lower electrical resistance of the interconnecting materials, has been noted by several authors [4–10]. Recently, Taur et al. stated that operation at lower temperature is seen as one of the available techniques to extend the improvement in performance described by Moore's law [11]. **Figure 1** shows Moore's law and several CMOS enhancements, including the use of low temperature.

©Copyright 2002 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/02/\$5.00 © 2002 IBM




Figure 1

Application of new structures, materials, and techniques to achieve continuation of Moore's law (dashed curve). Reproduced with permission from [3]. ILD = interlevel dielectric.

Early history—Electronic cooling prior to 1990

Many of the advances in computer technology which took place during the latter half of the twentieth century were made possible by both revolutionary and evolutionary increases in the packaging density of electronics. These advances began with the introduction of the transistor in 1947, and continue today with ultralarge-scale integration at the chip level coupled with the utilization of multichip modules. Over the past fifty years, increased circuit power dissipation and increased packaging density have led to substantial increases in chip and module heat flux, particularly in high-end computers. These trends are shown in **Figure 2** using data obtained from IBM and non-IBM products. During this time, the thermal design goal was to limit the magnitude of the rise in chip temperature above ambient temperature in order to ensure satisfactory electrical circuit operation and reliability.

Virtually all commercial computers were designed to operate at temperatures above ambient, generally in the range of 60 to 100°C. As circuit densities increased and heat removal became increasingly difficult, low-temperature enhancement of CMOS electrical performance was considered. A number of cooling technologies are available for providing cooling to low temperatures with a wide variety of available options. These cooling options and their corresponding temperature ranges are shown in **Figure 3**. In some cases only a few watts of heat can be removed by the cooling technology, whereas the portion of the computer system that requires a significant amount of cooling, primarily the

processor chips, can dissipate heat in the hundreds of watts range. With power levels of this magnitude, vapor-compression refrigeration appears to be one of the most applicable options.

Vapor-compression refrigeration was not always the primary cooling technique for lowering the temperature of electronic hardware. Starting in the mid-1970s, work began at IBM and elsewhere on a cryogenic computer based on superconducting logic. The necessary packaging technology and room-to-low-temperature interfaces were developed [12]. However, because of the rapid pace at which silicon technology progressed, both in terms of speed and level of integration, the performance margin of the superconducting logic computer was steadily eroded, and most of these programs were terminated. In late 1986, ETA Systems, Inc., a subsidiary of Control Data Corporation, shipped the first liquid-nitrogen-cooled CMOS computer system. This machine utilized direct-immersion cooling of single-chip modules immersed in liquid nitrogen [13]. Although this system utilized a unique low-temperature cooling technology, only a few of these systems were ever shipped. For most applications, "open cycle" cooling by consumption of a liquid cryogen was seen as impractical, leaving as the only alternative, at present, the mechanical refrigerator. Gaensslen et al. [14] noted that if low-temperature operation is ever to become a commercial reality for medium- and large-size digital computers, some form of closed-cycle refrigeration system will be required, because cryogenic fluids are too costly to be constantly expended in an open-pool cooling mode. Since the vapor-compression refrigeration system has been reliable and in use for many years, it is an obvious choice to consider for electronic cooling.

Vapor-compression refrigeration was proposed in 1805, and a working model was constructed around 1834 [15]. One of the earliest successful vapor-compression refrigeration machines, developed by Charles Tellier in France, was installed in a New Orleans brewery to preserve the flavor of beer and ale by keeping it at a constant temperature. In the U.S., demand for reliable cooling for brewing and later for ice-making led to the establishment of numerous firms that could design, manufacture, and install complete refrigeration systems. The refrigeration business was well established by 1900 [16].

In 1924, General Electric introduced the first domestic refrigerator with a hermetically sealed motor and compressor. Until that time, compressors were belt-driven, and shaft seals on the compressor were prone to leak refrigerant. The development of the automotive air conditioner began in earnest in 1930 when General Motors Research Laboratories conceived the idea of the vapor-compression system with R-12 refrigerant [17]. These vapor-compression refrigeration systems were the

advent of the modern refrigeration systems employed in many applications and were the primary reason for major growth in the South and Southwest over the last 50 years. As is evident, the technology is well established, but it has not been applied to cooling electronic hardware until very recently. Since refrigerators are widely available, and since the compressor and fan are the only moving parts in the cooling system, the building blocks embody a stable, reliable, and mature technology.

Recent history—Electronic cooling in the 1990s

One of the first companies to employ a vapor-compression refrigeration system in modern computers was Kryotech, Inc. The Kryotech technology has its roots in the National Cash Register (NCR)/Intel “Cheetah” project, which was an effort to boost the Intel Pentium** processor performance by cooling the chip to low temperature and boosting the CPU clock [18, 19]. NCR had a Cheetah unit ready to show at COMDEX Fall '94, but elected not to demonstrate or subsequently manufacture the product. Shortly thereafter, five longtime NCR employees obtained patent and technology licenses from their employer, formed Kryotech, and set forth to exploit the enhancement capabilities achieved with low-temperature refrigeration. Kryotech teamed with Digital Equipment Corporation to show a refrigerated workstation running at 767 MHz in 1996 [20].

With the increase in CMOS performance achieved with lower temperatures, a number of companies have embarked on programs to investigate cooling electronics, some evolving to major product announcements and shippable products. DEC, AMD, Sun Microsystems, SYS Technologies, and Kryotech, Inc. [21] have all shown computers that utilize the vapor-compression refrigeration cycle. In September 1997, IBM began shipping its largest S/390* servers with refrigeration, followed by the next three generations of IBM systems using refrigeration in 1998, 1999, and 2000 [22–24]. Finally, in December 1999, Kryotech teamed up with Advanced Micro Devices to market the first computer to achieve 1-GHz speed [21]. In achieving this speed, the evaporator attached to the Athlon** processor was cooled to -40°C with a vapor-compression refrigeration system [21]. All of these systems use a conventional refrigeration system to maintain chip temperatures below those of comparable air-cooled systems, but well above cryogenic temperatures.

Advantages of low-temperature electronic cooling

Performance

The primary motivation for utilizing lower temperature for cooling CMOS circuits has been the increased performance achieved with lower MOSFET junction




Figure 2

Trends in high-end computers: (a) Module heat flux; (b) chip power. (Courtesy of Stephen Kosonocky, IBM Research Division.)




Figure 3

Electronic cooling techniques and their corresponding operating temperature ranges.




Figure 4

Drain current, I_d , in a one-device memory cell. Reprinted with permission from [26]; ©1997 IEEE.

temperatures. In 1998, performance measurements of a CMOS single-chip module (SCM) showed that circuit speed increased at a rate of approximately 1.4% for every 10°C reduction in chip temperature (estimated from measurements of change in performance between 7°C and -20°C).¹ Electron and hole mobilities are the primary electrical properties that improve with the lowering of temperature. Transistor switching speed is proportional to the mean carrier (electron and hole) velocity in the device; and mobility is the ratio of the carrier (electron or hole) velocity to electric field. Mobility increases as temperature decreases because of the reduction of carrier scattering caused by thermal vibrations of the semiconductor crystal lattice. Future performance improvements will diminish as chip technologies approach a 100-nm gate length [25]. The advantage of a higher mobility is greatly diminished in such small devices because of increased scattering from the contacts and decreases in electron mobilities due to higher electric fields and carrier temperatures [6].

Leakage current reduction

Another important improvement with low temperature is in the reduction of leakage current. This can best be illustrated using a one-device dynamic memory cell [26]. The length of time during which information can be stored on the capacitor is limited by the p-n junction leakage, I_j , and the off-state device leakage, I_d . **Figure 4** shows the behavior of the drain current, I_d , as the temperature is reduced from 100°C to -50°C. The device (100°C curve) has a leakage current of 10^{-8} A when the gate voltage is

zero. If the temperature is reduced to -50°C, I_d decreases to 10^{-13} A for a gate voltage of zero. This is a decrease in leakage current of five orders of magnitude. The memory cell information retention time increases by the same factor. By reducing the temperature from 100°C to -50°C, the junction leakage, I_j , is reduced by many orders of magnitude, thereby eliminating device leakage as a problem.

Functionality

Increasing chip power may require low-temperature cooling to maintain chip junctions within their functional temperature limits. As chip power increases above 125–150 W [27], low-temperature cooling technologies may be the only means of maintaining chip temperatures within functional temperature limits. In addition, low-temperature operation enables a lower chip burn-in temperature, thus reducing chip leakage currents and minimizing chip over-current yield losses during the burn-in operation [28]. However, even though functionality may become the primary reason for the use of lower temperatures, the performance gains (higher speed and/or lower leakage currents) achieved will still be realized.

Reliability

Lowering the chip temperature is expected to improve the reliability of the overall system. Nearly all degradation mechanisms in electronic devices, such as interdiffusion, corrosion, and electromigration, have a thermal-activation component that follows the Arrhenius relationship. Since the rate of degradation decreases exponentially with decreasing temperature, orders of magnitude improvement in reliability could be expected with cooling. However, although the effect of temperature on reliability has been thoroughly verified empirically for elevated temperatures, it has yet to be demonstrated for reduced temperatures in densely populated chips. Detrimental effects of mechanical stress and strains which result from differences in thermal expansion of the various materials used in electronic components, coupled with local and overall temperature differences, arise when the system is cycled between room temperature and a low operating temperature. The magnitude of such effects on circuit performance and reliability must be considered when the temperature excursions are from room temperature to system operation below room temperature.

For example, tests described by Needham et al. [29] show that factors causing low-temperature-related failures occurred during the manufacturing of processors. In one test, 50 out of 20000 dies (chips) failed unique low-temperature tests at 0°C. Many of these fails were timing-related; that is, the device slowed as the temperature decreased. In another test, 53 dies passed ac testing at 25–100°C but

¹ A. Sutcliffe, IBM, private communication, 1998.

failed in tests performed at 0°C. Analysis of three of the fails showed that at lower temperatures a structure on the die opened electrically, causing a higher resistance and thereby failure. In two of the dies, the resistance in an electrical path increased to an unacceptable extent; in the third die, a high resistance in a via was observed. Unless these fails are carefully screened during manufacturing, failures could occur during field operation at low temperatures. Needham et al. reported that the defects, commonly known as soft defects, do not always cause failure at all conditions of temperature and voltage. They have shown that the percentage of soft fails to hard fails becomes greater as the complexities of future chip technologies are increased. This soft-failure mode increases the need to understand and eliminate reliability concerns in a low-temperature operation.

Disadvantages of low-temperature electronic cooling

Cost

Adding cooling hardware to the electronics adds cost to the system. However, this disadvantage, along with others such as space and power for the refrigeration hardware, must be weighed against the advantages of performance, leakage current, functionality, and reliability. The cost of cooling a workstation microprocessor with nonredundant cooling hardware, while not high in actual dollars, is a large fraction (10–20%) of the total cost of the system. For a large server the cost of cooling is higher in actual dollars, since cooling hardware redundancy is required (redundant cooling hardware continues cooling the system if the primary cooling hardware fails, thereby maintaining continuous operation of the system). However, the cooling hardware cost added in percentage of total system cost is small because of the high overall cost of the system. The space required for cooling hardware is large compared to the hardware being cooled, but the key factor to consider is whether the space occupied by the cooling hardware would be better used by additional, conventionally cooled microprocessor hardware. There is no advantage to adding refrigeration unless the chip modules cannot be maintained at acceptable temperatures with another cooling technique requiring a smaller volume to package, such as air cooling. On the other hand, if the system is not volume-constrained, adding refrigeration hardware may be very attractive.

Power

One other aspect of employing refrigeration within a computer system is the type of power required to operate the refrigeration hardware during start-up and during normal operation. The highest power is drawn for refrigeration hardware during compressor start-up.

Consideration should be given to whether ac or dc operation should be chosen. Alternating-current operation causes a higher peak current and therefore requires higher power during start-up than direct-current operation. The ac power draw during start-up is especially high if the compressor motor is single-phase. In some cases, the peak current can be five to seven times higher than the operating current, causing strain on the power system.

Application to S/390 high-end servers

In 1993 IBM began transforming its mainframe computer technology from bipolar emitter-coupled logic (ECL) to CMOS logic because of the tremendous advantage offered by CMOS in circuit density. Rather than requiring hundreds of ECL chips, a high-performance central processing unit may be contained on a single CMOS chip. Also, the low switching current of CMOS greatly reduces power consumption, eliminating the need for a complex and expensive water-cooled package. The IBM S/390 G4 CMOS system, first shipped in 1997, can contain 12 processing units, up to two levels of cache, and the bus-switching logic packaged in a single multichip module (MCM) on one processor board. This G4 system delivers performance comparable to that of an IBM 9021-711 bipolar system, in which the corresponding logic occupies 56 multichip modules on 14 boards. From the point of view of thermal technology, this was the first IBM system to employ refrigeration cooling. The decision to employ refrigeration cooling instead of other cooling options such as high-flow air cooling or various water-cooling schemes focused on the system performance improvement realized with the refrigeration system. The average processor temperature for G4 was approximately 40°C, a temperature improvement of 35°C compared to an air-cooled system of the same design. This 35°C lower operating temperature enabled the frequency (speed) of the processor chips to be increased.

The IBM G4 server is shown in **Figure 5** [30]. The bulk power subassembly for the system, shown at the top of the frame, distributes 350 Vdc throughout the frame. Below the bulk power is the central electronic complex (CEC) where the MCM, housing the 12 processors, is located. Various electronic “book” packages (memory, control modules, dc power supplies, etc.) are mounted on each side of the processor module. Below the CEC are blowers that provide air cooling for all of the components in the CEC except for the processor module, which is cooled by refrigeration. Below the blowers are two modular refrigeration units (MRUs) which provide cooling via the evaporator mounted on the processor module. Quick connects located at the evaporator permit concurrent maintenance of the MRU while the system is operating. In the bottom of the frame, the I/O electronic books are




Figure 5

S/390 G4 refrigerated system. Reprinted with permission from [30]; ©2000 ASME.




Figure 6

Side view of G4 frame showing airflow path.

installed, along with the associated blowers to provide the air cooling. In addition to cooling the I/O books, these blowers also provide cooling for the condensers within the MRUs. The airflow path is shown in a side view of the frame in **Figure 6**.

The refrigeration system was designed such that only one MRU is operated at a time during normal operation. Should one MRU fail, the other one turns on automatically. The failed MRU can be replaced via quick connects located at the evaporator. The evaporator mounted

on the processor module is fully redundant in that two independent loops utilizing copper tubes are interleaved through a thick copper plate. Refrigerant passing through one loop is adequate to cool the MCM under all environmental extremes allowed by the system.

The follow-on G5 and G6 systems were packaged in a similar fashion and with an almost identical refrigeration system. Some of the key MCM characteristics of the G4–G6 are highlighted in **Table 1**, and a cross section of the MCM is shown in **Figure 7**. With lower MCM powers for G5 and G6, a lower operating temperature was permitted with the same refrigeration hardware, but several design changes were implemented within the system to protect against condensation. These are discussed later in this paper.

The implementation of the refrigeration hardware to cool a multichip module required additional space for this hardware, because frame designs for CMOS processors up to this time did not provide sufficient space. Housing the MRUs in a separate frame was considered but was deemed unacceptable after concerns were raised about testing the separate frame, shipping it, then reconnecting the frame in the field. To accommodate the hardware within the frame, more space had to be provided. Early sizings on the hardware, showing that the MRU could fit within 6 U (1 U = 1.75 in.) increased the total height of the frame to 42 U. Although concern was expressed about the overall frame height for shipping, the power assembly located on top of the frame was removable if required for ease of installation. By installing the complete refrigeration system within the single frame, the system could be shipped and installed without breaking connections after completion of final testing within manufacturing.

Operation of the refrigeration hardware was a concern because the power available for the entire system was restricted by the space allotted for power supplies. One reason for using a dc compressor was to minimize the start-up operating current. An off-the-shelf single-phase ac compressor providing the same refrigeration capability would have been marginal or unacceptable from a power usage standpoint during start-up.

Refrigeration system thermal design

The design and selection of the components within the MRU and evaporator are described below, with emphasis placed on the final design. The selection of the refrigeration components to cool the 1050-W G4 multichip module involved an experimental program that was compressed into a short period of time because of program schedule constraints. A total of five condensers, four compressors, two filter-dryers, five expansion valves, two evaporators, two hot-gas bypass valves, and two accumulators were evaluated in various system configurations. With the complexity of a two-phase

cooling system, the choice of the components in many cases reflected the relationships among components, so various combinations of components were tested to achieve the best overall system performance from the standpoint of cost, performance, and reliability.

R-134a (CH_2FCF_3), a common refrigerant used in automotive and home air-conditioning systems, was selected as the refrigerant because of its environmental compatibility [31]. R-134a refrigerant does not contain chlorine and therefore does not contribute to the depletion of the ozone layer. The ozone depletion potential (ODP) of a material is a measure of its ability to destroy stratospheric ozone. Since R-134a contains no chlorine, its ODP is zero. Halocarbon refrigerants also contribute to global warming and are considered greenhouse gases. The global warming potential (GWP) of a greenhouse gas is an index describing its relative ability to trap radiant energy. The index is based on carbon dioxide, which has a very long atmospheric lifetime. R-134a does exhibit some GWP but is one of the preferred refrigerants.

All of the refrigeration system components except for the evaporator are housed within the MRU, shown in **Figure 8** with its cover removed. The MRU measures 267 mm $\times$ 267 mm $\times$ 711 mm and weighs approximately 27 kg. A schematic diagram of the refrigeration cycle, along with all of the components in the refrigeration system, is shown in **Figure 9**.

The compressor used in the MRU is a rotary compressor with a brushless dc motor providing speed control and thereby some control of the refrigeration system. This type of compressor uses a roller mounted on the eccentric of a shaft with a single vane or blade positioned in the nonrotating cylindrical housing. The blade reciprocates in a slot machined in the cylindrical block. Besides providing for a dc motor, the materials within the compressor were modified to provide long life for the refrigerant and oil used in the system.

High-pressure superheated vapor exits the rotary compressor and enters the air-cooled fin and tube condenser. As stated earlier, the exhaust air from the I/O cage located in the bottom half of the frame cools the condenser. After testing five different condensers, a




Figure 7

Cross section of MCM with evaporator.




Figure 8

Modular refrigeration unit (MRU).

three-row $\frac{3}{8}$ -in.-outer-diameter tube condenser with 14 aluminum fins per inch provided the desired performance. The heat is transferred from the refrigerant by 1) de-superheating, 2) condensing, and 3) subcooling. (The

Table 1 MCM comparisons for systems G4, G5, and G6.

System	Processor unit chip size (mm)	Processor unit chip power (W)	CP average temperature ($^{\circ}\text{C}$)	No. of processor unit chips	Total no. of chips	Total MCM power (W)
G4	17.3 $\times$ 17.3	45	35	12	29	1050
G5	15.5 $\times$ 15.5	20	20	12	29	650
G6	15.5 $\times$ 15.5	27	15	14	31	850




Figure 9

Schematic of refrigeration components. Reprinted with permission from [30]; © 2000 ASME.

term *superheating* refers to vapor at a temperature above saturation temperature; the term *subcooling* refers to a liquid at a temperature below saturation temperature.) Typically, the de-superheating and subcooling regions of the condenser are 5 to 10% of the overall condenser heat transfer area and are no different in the current MRU condenser. It was important in testing the five potential condensers in the MRU that 1) they provided sufficient subcooling that the inlet to the expansion device was liquid, and 2) the amount of airflow required through the fins of the condenser was minimized and compatible with that required for the I/O portion of the system. Furthermore, it was advantageous to limit the pressure drop through the refrigerant side of the condenser in order to improve the overall system performance. However, increasing subcooling too much was detrimental, since the pressure drop through the condenser began to increase dramatically. Consequently, the degree of subcooling permitted for the evaporator heat load had to be balanced against the pressure drop in the condenser. The amount of subcooling in this design varied from 2 to 5°C. On the basis of air-side temperatures, flow rates, and refrigerant temperatures, the refrigerant-side heat-transfer coefficient within the condenser ranged from 700 to 900 W/m²·K. This is similar to the data provided by

Outokumpu Copper Franklin, Inc. [32] for smooth tube condensation.

A filter downstream from the condenser removes any moisture, acids, oil decomposition products, insoluble materials such as metallic particles, and copper oxide or other contaminants in the refrigerant system. The amount of moisture in a refrigerant system must be kept below an allowable limit to provide satisfactory operation. Moisture is removed from components during manufacture and assembly to minimize the amount of moisture in the completed assembly. However, as an added precaution, a filter is installed downstream of the condenser to provide extra protection against moisture in the system. The filter-dryer is not a substitute for poor workmanship or design, but a maintenance tool necessary for continued and proper system performance.

The subcooled liquid leaving the condenser and filter-dryer enters the thermostatic-expansion (TX) valve, which controls its flow to the evaporator in response to the degree of superheating of the gas leaving the evaporator. The thermostatic-expansion valve functions to keep the evaporator active without permitting liquid to return through the suction line to the compressor. This is accomplished by controlling the mass flow of the refrigerant entering the evaporator so that the flow rate equals the rate at which the refrigerant can be completely vaporized in the evaporator by the absorption of heat from the microprocessors. Because the expansion valve is operated by superheating and responds to changes in superheating, a portion of the evaporator must be used to superheat the refrigerant gas.

It was essential that the degree of subcooling provided by the condenser be maintained to prevent a reduction in the capacity of the expansion valve caused by flashing of the liquid refrigerant. Thermostatic-expansion-valve ratings are based on vapor-free liquid entering the valve. If gas is present in the entering liquid, the valve capacity is reduced substantially because the gas must be handled along with the liquid. Flashing of the liquid refrigerant may be caused by a pressure drop in the liquid line or the filter-dryer, or a combination of these. A thermostatic-expansion valve must provide stable evaporator temperatures during normal operation, and it must not allow excessive evaporator temperature oscillations during start-up. Only a few expansion valves satisfied both of these temperature requirements. Extreme evaporator temperature oscillations occurred during start-up with some expansion valves. These oscillations could be attributed to a phenomenon referred to in the industry as "hunting"—an alternate overfeeding and starving of the refrigerant feeding the evaporator. Although hunting is commonly attributed to the thermostatic-expansion valve, the latter is seldom solely responsible for this phenomenon. One reason for hunting is that all

evaporators have a time lag. When the thermostatic-expansion valve bulb signals for a change in refrigerant flow, the refrigerant must traverse the entire length of the evaporator channel before a new signal reaches the bulb. This lag or time lapse may cause constant overshooting of the valve during both opening and closing. Extreme hunting reduces the capacity of the refrigeration system because the mean evaporator pressure and temperature are lowered and the compressor capacity is reduced. After numerous tests, a combination of valve superheating, correct thermostatic element, bulb location, and suction piping size was required to provide a stable operating system.

The low-quality liquid (mostly liquid with very little vapor) leaving the expansion device travels through an insulated hose to quick connects located at the evaporator. The quick connects allow the MRU to be removed without affecting the evaporator or cooling of the MCM. The refrigerant flows in the evaporator in a series of two rows of $\frac{3}{8}$ -in. internally enhanced tubes with internal ridges. The internal ridges improve the heat-transfer characteristics of the copper tubes soldered into the evaporator copper block. The heat-transfer coefficient based on the internal diameter of the tube ranged from 1400 to 1600 W/m²·K, similar to the values reported by Zurcher et al. [33]. The tube paths for each MRU are intertwined within the copper block so that either MRU loop can support the full heat load of the MCM. The thermodynamic states for the refrigeration cycle just described are shown in **Figure 10** for the G4 system.

It was essential that the temperature of the MCM be maintained relatively constant, regardless of changes in the ambient temperature, altitude, and heat load of the I/O cage. For these reasons, a hot-gas bypass valve was added to maintain a steady evaporator temperature and thus a relatively constant MCM temperature. The hot-gas bypass is pulsed on the basis of the MCM temperature monitored by thermistors mounted in the evaporator. Too low an evaporator temperature causes the hot-gas bypass valve to pulse, while too high a temperature results in the valve not pulsing at all. The hot-gas bypass valve within the system arrangement is shown in Figure 9. The bypass valve is connected between the low-pressure side of the expansion valve and the entrance to the condenser, an arrangement which prevents overheating of the compressor if it is used for protracted periods of time. Other arrangements of the hot-gas bypass valve can cause serious problems with the operation of the compressor. The hot-gas bypass lines were selected so that pressure loss is only a small percentage of the pressure drop across the valve. In addition to maintaining a tight temperature control of the evaporator, the hot-gas bypass valve was necessary to 1) unload the compressor to reduce starting




Figure 10

S/390 refrigeration cycle (R-134a). Reprinted with permission from [34]; © 1998 ASHRAE.

torque requirements and 2) permit capacity control down to 0% load conditions without stopping the compressor.

The amount of refrigerant needed for proper system operation under all environmental conditions and for maximum protection against refrigerant leaks was key in the development phase of the program. Extremes in operating conditions were tested with different fill volumes to determine the maximum allowable while still permitting efficient operation of the refrigeration system. Too much refrigerant would cause a degradation of the condenser, while too little would result in too much superheating at the evaporator. After selection of the components that provided the lowest cost and most efficient cycle along with the maximum refrigerant charge, the coefficient of performance (COP) for various environmental extremes was measured. For the G4 system, the COP varied from 2 to 3 depending on the environmental conditions.

Condensation control

To avoid moisture condensation on the MCM hardware [35], especially the ceramic substrate and the connector pins, all of the cooling hardware including the evaporator copper cold plate is contained in an airtight metal enclosure (**Figure 11**) with one open face. The open face has a specially designed gasket that seals the evaporator enclosure to the planar circuit board (the planar board is not shown in Figure 11). A cavity inside the evaporator enclosure is designed to house approximately 260 grams of silica gel desiccant, an amount deemed sufficient to absorb any moisture leaking into the enclosure over the life of the machine.

To test various seal materials and to simulate the actual system configuration, the planar board was replaced by a




Figure 11

Evaporator seal humidity test setup.




Figure 12

Relative humidity values recorded inside the evaporator cavities using test setup of Figure 11.

flat plate with two humidity sensors placed on the top portion and the bottom portion of the plate, as shown in Figure 11. The evaporator assemblies under test contained 26 grams of silica gel, one-tenth the amount used in the actual application. Testing was performed by placing the evaporator enclosures in a humidity chamber set at 50°C and 80% relative humidity. The relative humidity recorded is shown in **Figure 12** as a function of time. After two months of testing, the evaporator with the butyl rubber gasket showed a negligible moisture ingress rate, and the increase in relative humidity was less than 1%. The butyl rubber gasket was selected because it provided significantly

better moisture sealing capability than a natural rubber gasket. To further prevent condensation, the evaporator is encased in insulation and sealed with an outer metal cover to prevent any possibility of moisture infusion and subsequent condensation in the area around the module.

Because the heat dissipated by the MCMs in G5 and G6 was less than that in G4, lower evaporator temperatures were implemented for the same refrigeration system. However, because of the lower temperatures and to provide additional moisture control, a backup desiccant system was used. To prevent moisture from leaking into the evaporator enclosure, the enclosure air pressure was raised by 0.1 psi above the atmospheric pressure, and, with the help of the cage blowers, air was forced into a small copper tube connected to the enclosure. On its way to the enclosure, this forced airstream was dried by passing it through a canister filled with desiccant. The chosen desiccant changes color when it has absorbed too much moisture, and the canister containing it is made of a transparent material; thus, the remaining life of the desiccant can easily be observed. Only a very minuscule amount of air flows through this backup desiccant system because the evaporator enclosure is essentially airtight. The backup desiccant should therefore last for the life of the system. To verify that the seal between the planar board and evaporator enclosure remains adequately airtight during manufacturing, a slight vacuum is applied to the enclosure. If the vacuum holds for a certain period of time, the system is deemed sufficiently airtight to protect against moisture ingress over the life of the machine.

Mechanical design

To ensure a highly reliable cooling system, the MRU components were qualified using a series of tests and thorough reviews of available test and field data from the component suppliers. Ten MRUs were subjected to life test for a period of two years to determine the amount of wear on the compressor bearings and the pistons. In addition, the wear measurement of the hot-gas bypass valves was conducted after six million cycles to evaluate the extent of wear on the valve push rod. The results showed that while there was a slight amount of wear in the compressor and the hot-gas bypass valve components, there was no degradation in the performance that would affect the performance or reliability of the system.

In addition to the compressor and hot-gas bypass wearout mechanisms, the prototype MRU copper pipes suffered from premature fatigue cracking. The prototype design failed in an average of less than 30 hours of operation when the compressor was run at 1800 rpm. On the basis of experience with copper tube cracking in prototype MRUs, work was initiated to modify the piping design assemblies and, more significantly, to quantitatively

measure improvements in the fatigue life (number of cycles to failure). The cyclic strain at the failing locations was measured as a function of compressor frequency. It was found that the cyclic strain peak amplitude at 1800 rpm was in the range 367–640 micro-mm/mm. The cyclic strain peak amplitude at 2300 rpm was in the range of 99–432 micro-mm/mm, or microstrain. For the copper piping with a tensile modulus of elasticity of 117 GPa [36], the stress is 75 GPa for 640 micro-mm/mm strain amplitude and 52 GPa for 432 micro-mm/mm strain amplitude. The fatigue strength (after 20 million cycles of the repeated application of the load) is 75 GPa [37]. Thus, at 1800 rpm the piping is anticipated to fail at 20 million cycles or 7.7 days. The stress at 2300 rpm is 33% lower than the fatigue strength. This reduced strain at 2300 rpm should correspond to fewer tube failures due to fatigue, which is confirmed by ten MRU life test results. The piping design was changed by introducing more loops, especially in the inlet and outlet pipes attached to the vibrating compressor. As shown in Figure 8, the piping around the compressor has several loops to reduce energy transmission to joints and other parts of the pipes. To determine the areas of high strain in the piping system, the MRU operated at various frequencies from 1800 to 5000 rpm. At each 100-rpm increment, a stroboscope and a high-speed video camera were used to locate regions that exhibited high displacements, indicative of high strains. These high-strain locations were typically located around piping terminations, such as the piping connected to the condenser, compressor inlet and outlet connections, the small and large accumulator inlet and outlet connections, hot-gas bypass connections, and t-joint connections. The cyclic strain amplitudes were measured in the high-strain areas. A graph of cyclic strain amplitudes with respect to the small accumulator elbow, a region of high strain as measured by the stroboscope, is shown in **Figure 13**. The strain amplitude at a frequency of 2300 rpm was less than 100 micro-mm/mm, which is significantly less than when the compressor was run at 1800 rpm. The new design piping is predicted to be free of fatigue cracks. Various compressor speeds, along with pulsing of the hot-gas bypass valve, are needed in order to maintain tight temperature control of the processor module. By pulsing every 20 seconds, six million cycles is comparable to 3.8 years in the field.

In addition to the cyclic strain amplitude measurements described above, a photostress analysis [37] was utilized to highlight the location and magnitude of the high cyclic strain. In order to conduct the photostress analysis, the piping was coated with a special strain-sensitive coating. When the coated piping was illuminated with a polarized light, the coating showed a fringe pattern of piping strain. By comparing the fringe pattern to a standard fringe




Figure 13

Typical piping strain versus rpm for two different modular refrigeration units (microstrain = 10^{-6} strain).

pattern, an estimate of the strain magnitude was made for high dynamic strain around the piping terminations. This technique is also used to reveal residual stresses. When the coated unit is subjected to dynamic strain during shipping, a high residual strain is identified in the coating as a fringe.

The maximum dynamic strain recorded during operation was identified at 700 micro-mm/mm. In addition, a high residual strain (1000 micro-mm/mm) on the piping near the condenser and large accumulator, due to insufficient packaging of the MRU during shipment, was identified. After improvement of the MRU shipping packaging, this residual stress was significantly reduced below 80 micro-mm/mm.

Summary

A refrigeration system has been designed and used in the IBM S/390 G4 large-scale server. Similar cooling systems have been used in the subsequent-generation systems, G5, G6, and z900, each achieving lower processor temperatures. The cooling system design included careful selection of the refrigeration components, condensation control through the use of desiccants and strategically located insulation, and temperature control primarily through the use of a hot-gas bypass valve. Additionally, design of the piping and control of compressor speed produced a design free of fatigue fails within the life of the system. Many components were evaluated in the search for the lowest-cost, most reliable, and most thermodynamically efficient design. Various refrigeration-system issues (condensation control, cracking of piping, thermal performance, etc.) were resolved during the design cycle. Applied to IBM S/390 systems, this cooling technology has proved to be a viable technique for use in large-scale servers.

Acknowledgments

The authors wish to acknowledge Udo Jourdan, William Winkler, James Gutelius, Dave McClafferty, and Mark Marnell for their assistance in collecting the test data, and P. J. Singh for his helpful comments. The authors gratefully acknowledge the continuous encouragement and support offered by Vincent Cozzolino and Nancy Drumm.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Intel Corporation or Advanced Micro Devices Corporation.

References

- G. E. Moore, "Cramming More Components onto Integrated Circuits," *Electron.* **38**, 114–117 (1965).
- G. E. Moore, "Progress in Digital Integrated Electronics," *IEEE International Electron Devices Meeting (IEDM) Digest of Technical Papers*, December 1975, pp. 11–13.
- R. D. Isaac, "The Future of CMOS Technology," *IBM J. Res. & Dev.* **44**, No. 3, 369–378 (2000).
- Low-Temperature Electronics*, R. K. Kirschman, Ed., IEEE Press, New York, 1985.
- R. K. Kirschman, "Cold Electronics: An Overview," *Cryogenics* **25**, No. 3, 115–122 (1985).
- P. M. Solomon, "Materials, Devices, and Systems," *Low-Temperature Electronics*, R. K. Kirschman, Ed., IEEE Press, New York, 1985, pp. 16–19.
- R. W. Keyes, E. P. Harris, and K. L. Konnerth, "The Role of Low Temperature in the Operation of Logic Circuitry," *Proc. IEEE* **58**, 1914–1932 (1970).
- R. C. Jaeger and F. H. Gaensslen, "MOS Devices and Switching Behavior," *Low-Temperature Electronics*, R. K. Kirschman, Ed., IEEE Press, New York, 1985, pp. 90–93.
- F. H. Gaensslen and R. C. Jaeger, "Temperature Dependent Threshold Behavior of Depletion Mode Mosfets," *Solid-State Electron* **22**, No. 4, 423–430 (1979).
- R. C. Jaeger, "Development of Low Temperature CMOS for High Performance Computer Systems," *Proceedings of the IEEE International Conference on Computer Design: VLSI in Computers*, 1986, pp. 128–130.
- Y. Taur, D. A. Buchanan, W. Chen, D. J. Frank, K. E. Ismail, S. Lo, G. A. Sai-Halasz, R. G. Viswanathan, H. C. Wann, S. J. Wind, and H. Wong, "CMOS Scaling into the Nanometer Regime," *Proc. IEEE* **85**, No. 4, 486–504 (1997).
- M. B. Ketchen, D. J. Herrell, and C. J. Anderson, "Josephson Cross-Section Model Experiment," *J. Appl. Phys.* **57**, No. 7, 2550–2574 (1985).
- R. J. Krane, J. R. Parsons, and A. Bar-Cohen, "Design of a Candidate Thermal Control System for a Cryogenically Cooled Computer," *IEEE Trans. Components, Hybrids, Manuf. Technol.* **11**, No. 4, 545–566 (1985).
- F. H. Gaensslen, V. L. Rideout, E. J. Walker, and J. J. Walker, "Very Small MOSFETs for Low Temperature Operation," *IEEE Trans. Electron Devices* **ED-24**, 218 (1977).
- B. Nagengast, "Early Twentieth Century Air-Conditioning Engineering," *ASHRAE J.*, pp. 55–62 (1999).
- Anonymous, "Historical Review of the Rise of Mechanical Refrigeration," *Ice and Refrigeration*, pp. 46–58 (1901).
- M. S. Bhatti, "Riding in Comfort: Part II," *ASHRAE J.*, pp. 44–50 (1999).
- http://www2.sharkyextreme.com/hardware/kryotech_k63.
- A. Bindra, "Kryotech Takes NCR Cooling Schemes to Market," *Electron. Eng. Times*, p. 78, May 20, 1996.
- Anonymous, "Kryotech Hotrods Alpha to 767MHz, Creates Record 1.5 Gigaflips NT Workstation," *Client Server News*, No. 178, p. 1 (November 25–29, 1996).
- Kryotech News; see <http://www.kryotech.com/index2.html>.
- IBM S/390 G5 announcement, 1998; see <http://www-1.ibm.com/servers/eserver/zseries/>.
- P. Singh, W. D. Becker, V. Cozzolino, M. J. Ellsworth, Jr., R. R. Schmidt, and E. J. Seminaro, "System Packaging of a CMOS Mainframe," *Adv. Microelectron.* **25**, No. 7, 14–19 (1998).
- IBM S/390 G7 announcement, 1998; see [http://www2.ibm.link.ibm.com/cgi-bin/master?xh=iEtzPO5IG\\$CiO\\$1USenGnN9332&request=announcements&parms=H%5f102%2d052&xfr=N](http://www2.ibm.link.ibm.com/cgi-bin/master?xh=iEtzPO5IGCiO1USenGnN9332&request=announcements&parms=H%5f102%2d052&xfr=N).
- G. A. Sai-Halasz, M. R. Wordeman, D. P. Kern, S. A. Rishton, E. Ganin, T. H. P. Chang, and R. H. Dennard, "Experimental Technology and Performance of 0.1- μ m-Gate-Length FETs Operated at Liquid-Nitrogen Temperatures," *IBM J. Res. & Dev.* **34**, 452–465 (1990).
- A. R. Keshavarzi and C. F. Hawkins, "Intrinsic Leakage in Low Power Deep Submicron CMOS ICs," *Proceedings of the IEEE International Test Conference*, 1997, pp. 146–155.
- International Technology Roadmap for Semiconductor (ITRS)*, Semiconductor Industry Association (SIA); see <http://public.itrs.net/Files/2001ITRS/Home.htm>.
- B. W. Curran, Y. H. Chan, P. T. Wu, P. J. Camporese, G. A. Northrop, R. F. Hatch, L. B. Lacey, J. P. Eckhardt, D. T. Hui, and H. H. Smith, "IBM eServer z900 High Frequency Microprocessor Technology, Circuits, and Design Methodology," *IBM J. Res. & Dev.* **46**, No. 4/5, 631–644 (2002).
- W. Needham, C. Prunty, and E. H. Yeoh, "High Volume Microprocessor Test Escapes, An Analysis of Defects Our Tests Are Missing," *Proceedings of the IEEE International Test Conference*, 1998, pp. 25–34.
- R. R. Schmidt, M. J. Ellsworth, Jr., R. C. Chu, and D. Agonafer, "Refrigeration Cooled Computers: Application and Review," *Proceedings of the ASME International Engineering Congress and Exposition*, 2000, Vol. 28, pp. 277–283.
- A. D. Althouse, C. H. Turnquist, and A. F. Bracciano, *Modern Refrigeration and Air Conditioning*, The Goodheart-Willcox Company, Inc., Tinley Park, IL, 1996, p. 343.
- Technical Bulletin ACRH-1*, Outokumpu Copper Franklin, Inc., Franklin, KY, 1996.
- O. Zurcher, J. R. Thome, and D. Favrat, "Flow Boiling and Pressure Drop Measurements for R-134a/Oil Mixtures. Part 2: Evaporation in a Plain Tube," *Int. J. Heating, Ventilating, Air-Cond. & Refrig. Res.* **3**, No. 1, 54–64 (1997).
- R. H. Howell, H. J. Sauer, Jr., and W. J. Coad, *Principles of Heating, Ventilating, and Air Conditioning*, American Society of Heating, Refrigerating, and Air Conditioning Engineers, Inc., Atlanta, GA, 1998, p. 2.14.
- M. J. Ellsworth, Jr., R. R. Schmidt, and D. Agonafer, "Design and Analysis of a Scheme to Mitigate Condensation on an Assembly Used to Cool a Processor Module," *IBM J. Res. & Dev.* **46**, No. 6, 753–761 (2002, this issue).
- Copper Development Associates, Inc., New York, NY, *Standards Handbook: Copper, Brass, Bronze*, 1973, pp. 35–36.
- "Introduction to Stress Analysis by PhotoStress Method," *Technical Note TN 702-1*, Measurements Group, Inc., Raleigh, NC, 1992.

Received June 13, 2000; accepted for publication May 9, 2002

Roger R. Schmidt IBM Server Group, 2455 South Road, Poughkeepsie, New York 12601 (c28rrs@us.ibm.com). Dr. Schmidt is a Distinguished Engineer, a member of the IBM Academy of Technology, and an ASME Fellow. He received his B.S. degree in mechanical engineering from Bradley University and his M.S. and Ph.D. degrees in mechanical engineering from the University of Minnesota. Dr. Schmidt has more than 25 years' experience in engineering and engineering management in the thermal design of IBM large-scale computers. He has led development teams in cooling mainframes, client/servers, parallel processors and test equipment utilizing such cooling mediums as air, water, and refrigerants. Dr. Schmidt has published more than 50 technical papers, and he holds 33 patents in the area of electronic cooling. He is a member of the ASME Heat Transfer Division and an active member of the ASME K-16 Electronic Cooling Committee, and is an Associate Editor of the *Journal of Electronic Packaging*. Over the past 15 years, Dr. Schmidt has taught many mechanical engineering courses for prospective Professional Engineers; he has given seminars on electronic cooling at several universities.

Budy D. Notohardjono IBM Server Group, 2455 South Road, Poughkeepsie, New York 12601 (budy@us.ibm.com). Dr. Notohardjono is a Senior Engineer in the Mechanical Analysis Shock and Vibration Department in Poughkeepsie. He received his Ph.D. degree in mechanical engineering from the University of Wisconsin at Madison in 1984, and his M.B.A. in finance from New York University in 1997. He has received two IBM Outstanding Technical Achievement Awards and holds four U.S. patents. Dr. Notohardjono is a registered Professional Engineer in the state of New York.